



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/800,109	03/12/2004	Christina Woody Mercier	112056-0427	8926
24267	7590	08/10/2007	EXAMINER	
CESARI AND MCKENNA, LLP 88 BLACK FALCON AVENUE BOSTON, MA 02210			WASSUM, LUKE S	
		ART UNIT	PAPER NUMBER	
		2167		
		MAIL DATE	DELIVERY MODE	
		08/10/2007	PAPER	

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Office Action Summary	Application No.	Applicant(s)
	10/800,109	MERCIER ET AL.
	Examiner	Art Unit
	Luke S. Wassum	2167

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --
Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

1) Responsive to communication(s) filed on 30 May 2007.
 2a) This action is **FINAL**. 2b) This action is non-final.
 3) Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

4) Claim(s) 23-46 is/are pending in the application.
 4a) Of the above claim(s) _____ is/are withdrawn from consideration.
 5) Claim(s) _____ is/are allowed.
 6) Claim(s) 23-46 is/are rejected.
 7) Claim(s) _____ is/are objected to.
 8) Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

9) The specification is objected to by the Examiner.
 10) The drawing(s) filed on 12 March 2004 is/are: a) accepted or b) objected to by the Examiner.
 Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).
 Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).
 11) The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

12) Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
 a) All b) Some * c) None of:
 1. Certified copies of the priority documents have been received.
 2. Certified copies of the priority documents have been received in Application No. _____.
 3. Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

1) <input type="checkbox"/> Notice of References Cited (PTO-892)	4) <input type="checkbox"/> Interview Summary (PTO-413) Paper No(s)/Mail Date. _____.
2) <input type="checkbox"/> Notice of Draftsperson's Patent Drawing Review (PTO-948)	5) <input type="checkbox"/> Notice of Informal Patent Application
3) <input type="checkbox"/> Information Disclosure Statement(s) (PTO/SB/08) Paper No(s)/Mail Date _____.	6) <input type="checkbox"/> Other: _____.

DETAILED ACTION

Continued Examination Under 37 CFR 1.114

1. A request for continued examination under 37 CFR 1.114, including the fee set forth in 37 CFR 1.17(e), was filed in this application after final rejection. Since this application is eligible for continued examination under 37 CFR 1.114, and the fee set forth in 37 CFR 1.17(e) has been timely paid, the finality of the previous Office action has been withdrawn pursuant to 37 CFR 1.114. Applicant's submission filed on 30 May 2007 has been entered.

Response to Amendment

2. The Applicants' amendment, filed 30 May 2007, has been received, entered into the record, and considered.

3. As a result of the amendment, new claim 46 has been entered. Claims 1-22 have been previously canceled. Claims 23-46 are now pending in the application.

The Invention

4. The claimed invention is an apparatus providing coherent data copying operations where data replication is controlled by a source storage controller directly to a destination controller and managed by a remote application.

Priority

5. The examiner acknowledges the Applicants' claim to domestic priority under 35 U.S.C. § 120, as a continuation of application 09/375,819, filed 16 August 1999.

Claim Rejections - 35 USC § 103

6. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negatived by the manner in which the invention was made.

7. The factual inquiries set forth in *Graham v. John Deere Co.*, 383 U.S. 1, 148 USPQ 459 (1966), that are applied for establishing a background for determining obviousness under 35 U.S.C. 103(a) are summarized as follows:

1. Determining the scope and contents of the prior art.
2. Ascertaining the differences between the prior art and the claims at issue.
3. Resolving the level of ordinary skill in the pertinent art.
4. Considering objective evidence present in the application indicating obviousness or nonobviousness.

8. This application currently names joint inventors. In considering patentability of the claims under 35 U.S.C. 103(a), the examiner presumes that the subject matter of the various claims was commonly owned at the time any inventions covered therein were made absent any evidence to the contrary. Applicant is advised of the obligation under 37 CFR 1.56 to point out the inventor and invention dates of each claim that was not commonly owned at the time a later invention was made in order for the examiner to consider the applicability of 35 U.S.C. 103(c) and potential 35 U.S.C. 102(e), (f) or (g) prior art under 35 U.S.C. 103(a).

9. Claims 23, 25, 26, 28, 29, 31, 33, 34, 36, 37, 39-41 and 45 are rejected under 35 U.S.C. 103(a) as being unpatentable over **Meyer** (U.S. Patent 5,867,733) in view of **Ohran et al.** (U.S. Patent 5,649,152).

10. Regarding claim 23, Meyer teaches a storage device controller substantially as claimed, comprising:

- a) copy logic (see disclosure that the data storage controller transfers data directly between the first and second data storage devices under control of the data storage device controller, without employing the memory array and computer bus, said transferring constituting copying, rendering the claimed copy logic inherent, col. 4, lines 28-44; see also col. 2, lines 7-19);
- b) the controller being operable to receive a copy command specifying the source volume and a target volume (see disclosure that the data storage controller transfers data directly between the first and second data storage devices under control of the data storage device controller, without employing the memory array and computer bus, said transferring constituting copying, rendering the claimed copy command specifying the source volume and a target volume inherent, col. 4, lines 28-44; see also col. 2, lines 7-19);
- c) the controller being operable to receive a write command specifying the source volume (see col. 5, lines 48-60); and
- d) the copy logic being operable in response to receiving the copy command to generate and send one or more storage device commands to one or more storage devices for the source and target volumes to copy data from the

source volume directly to the target volume without having a file server in the data path (see disclosure that the data storage controller transfers data directly between the first and second data storage devices under control of the data storage device controller, without employing the memory array and computer bus, said transferring constituting copying, rendering the claimed copy logic inherent, col. 4, lines 28-44; see also col. 2, lines 7-19).

Meyer does not explicitly teach a storage device controller including snapshot logic.

Ohran et al., however, teaches a system for providing and maintaining snapshots, including:

- a) snapshot logic (see Abstract, disclosing that the reference is a method for providing a static snapshot; see also col. 1, lines 15-18);
- b) an internal cache (see disclosure of block association memory, element 108 of Figure 1; see also col. 4, lines 51-56, disclosing that the block association memory may be a portion of the RAM of digital computer 102);
- c) the system being operable to communicate with a replication manager to receive a snapshot command issued by the replication manager, the

snapshot command specifying a range of data bytes of a source volume (see disclosure of a user indicating that a static image [i.e., snapshot] of the mass storage system is desired, said indication being analogous to the claimed snapshot command, col. 4, lines 14-24; note also the disclosure that mass storage system 104 can be any writable block-addressable storage system, such as one or more disks or a partition of a disk, a partition being a fixed portion of a disk, col. 3, lines 50-56; see also col. 5, lines 23-41);

d) the snapshot logic being operable, in response to the snapshot command, to take a snapshot of the range, the snapshot including a snapshot map and snapshot data, the snapshot map being stored by the snapshot logic in the internal cache and the snapshot data being stored by the snapshot logic in a snapshot volume (see col. 4, lines 20-35; see also disclosure that preservation memory [i.e. the snapshot data] can be an area of memory, one or more disks, a partition of a disk, or a file stored on a disk, col. 3, line 66 through col. 4, line 1; see also disclosure of block association memory [i.e. the snapshot map] that is used to associate blocks stored in preservation memory with the unique addresses of blocks on the mass storage system, col. 4, lines 51-62); and

e) wherein the snapshot map and snapshot data are used to maintain coherency of any data that is requested (see disclosure that copies of blocks on the mass storage system are placed in preservation memory whenever they are going to be changed by a write operation, unless an entry for that block is already in the preservation memory, and furthermore that when a read is requested, the preservation memory is first checked to see if it contains a copy of the block, and that if the copy exists in preservation memory, that copy is returned, otherwise the block is read from the mass storage system, col. 2, lines 55-64; see also col. 4, lines 30-48 et seq.; see also drawing Figure 2).

It would have been obvious to one of ordinary skill in the art at the time of the invention to incorporate snapshot logic of **Ohran et al.** into the storage device controller of **Meyer**, since this would allow the system to create periodic backups for recovery in the event of a failure of the mass storage system, while ensuring that said periodic backup would not be rendered inconsistent in the case where said mass storage system was being updated by other programs as the backup copy is being made, col. 1, lines 20-50).

11. Regarding claim 31, Meyer teaches a method substantially as claimed, comprising:

- a) a storage device controller (see disclosure that the data storage controller transfers data directly between the first and second data storage devices under control of the data storage device controller, without employing the memory array and computer bus, said transferring constituting copying, rendering the claimed copy logic inherent, col. 4, lines 28-44; see also col. 2, lines 7-19);
- b) receiving at the storage device controller a copy command specifying a copy operation from a source volume and a target volume (see disclosure that the data storage controller transfers data directly between the first and second data storage devices under control of the data storage device controller, without employing the memory array and computer bus, said transferring constituting copying, rendering the claimed copy command specifying the source volume and a target volume inherent, col. 4, lines 28-44; see also col. 2, lines 7-19);
- d) in response to receiving the copy command, the storage device controller generating and sending one or more storage device commands to one or

more storage devices of the source and target volumes to copy data from the source volume directly to the target volume without having a file server in the data path (see disclosure that the data storage controller transfers data directly between the first and second data storage devices under control of the data storage device controller, without employing the memory array and computer bus, said transferring constituting copying, rendering the claimed copy logic inherent, col. 4, lines 28-44; see also col. 2, lines 7-19).

Meyer does not explicitly teach a storage device controller including snapshot logic.

Ohran et al., however, teaches a method for providing and maintaining snapshots, including:

a) receiving a snapshot command issued by the replication manager, the snapshot command specifying a range of data bytes of a source volume (see disclosure of a user indicating that a static image of the mass storage system is desired, said indication being analogous to the claimed snapshot command, col. 4, lines 14-24; note also the disclosure that mass storage

system 104 can be any writable block-addressable storage system, such as one or more disks or a partition of a disk, a partition being a fixed portion of a disk, col. 3, lines 50-56; see also col. 5, lines 23-41);

b) in response to the snapshot command, taking a snapshot of the range, the snapshot including a snapshot map and snapshot data, the snapshot map being stored by the snapshot logic in an internal cache and the snapshot data being stored by the snapshot logic in a snapshot volume (see col. 4, lines 20-35; see also disclosure that preservation memory [i.e. the snapshot data] can be an area of memory, one or more disks, a partition of a disk, or a file stored on a disk, col. 3, line 66 through col. 4, line 1; see also disclosure of block association memory [i.e. the snapshot map] that is used to associate blocks stored in preservation memory with the unique addresses of blocks on the mass storage system, col. 4, lines 51-62); and

c) wherein the snapshot map and snapshot data are used to maintain coherency of any data that is requested (see disclosure that copies of blocks on the mass storage system are placed in preservation memory whenever they are going to be changed by a write operation, unless an entry for that block is already in the preservation memory, and furthermore that when a read is requested, the preservation memory is first checked to see if it contains a

copy of the block, and that if the copy exists in preservation memory, that copy is returned, otherwise the block is read from the mass storage system, col. 2, lines 55-64; see also col. 4, lines 30-48 et seq.; see also drawing Figure 2).

It would have been obvious to one of ordinary skill in the art at the time of the invention to incorporate snapshot logic of **Ohran et al.** into the storage device controller of **Meyer**, since this would allow the system to create periodic backups for recovery in the event of a failure of the mass storage system, while ensuring that said periodic backup would not be rendered inconsistent in the case where said mass storage system was being updated by other programs as the backup copy is being made, col. 1, lines 20-50).

12. Regarding claim 39, **Meyer** teaches a computer-implemented method substantially as claimed, comprising:

a) using a replication manager to manage a source storage device controller and a destination storage device controller, the source storage device controller being operable to control access to a source data object and the destination

device controller being operable to control access to a destination data block, the storage device controllers being operable to issue storage device commands (see disclosure that the data storage controller transfers data directly between the first and second data storage devices under control of the data storage device controller, without employing the memory array and computer bus, col. 4, lines 28-44; see also col. 2, lines 7-19; see also disclosure that the data storage device controller includes first and second device controllers coupled to the first and second storage devices, respectively, col. 4, lines 19-21); and

b) copying each block of the source data object to a corresponding block in the destination data object wherein the data is directly transferred between the source and destination storage device controllers without traversing a server operable to process file system requests (see disclosure that the data storage controller transfers data directly between the first and second data storage devices under control of the data storage device controller, without employing the memory array and computer bus, said transferring constituting copying, col. 4, lines 28-44; see also col. 2, lines 7-19; see also disclosure that the data storage device controller includes first and second

device controllers coupled to the first and second storage devices, respectively, col. 4, lines 19-21).

Meyer does not explicitly teach a system including snapshot logic.

Ohran et al., however, teaches a system for providing and maintaining snapshots, including:

a) internally generating within the source storage device controller in communication with the replication manager, a snapshot version for each block of the source data object changed by one or more write operations to the block during the course of a copy operation (see col. 4, lines 20-35; see also disclosure that preservation memory [i.e. the snapshot data] can be an area of memory, one or more disks, a partition of a disk, or a file stored on a disk, col. 3, line 66 through col. 4, line 1; see also disclosure of block association memory [i.e. the snapshot map] that is used to associate blocks stored in preservation memory with the unique addresses of blocks on the mass storage system, col. 4, lines 51-62; see also disclosure that copies of blocks on the mass storage system are placed in preservation memory whenever they are going to be changed by a write operation, unless an

entry for that block is already in the preservation memory, col. 2, lines 55-

64; see also col. 4, lines 30-48 et seq.; see also drawing Figure 2);

b) copying each block of the source data object to a corresponding block in the destination data object in the absence of the snapshot version of the block and otherwise copying the snapshot version of the source data object block to the corresponding block in the destination data object (see disclosure that when a read is requested, the preservation memory is first checked to see if it contains a copy of the block, and that if the copy exists in preservation memory, that copy is returned, otherwise the block is read from the mass storage system, col. 2, lines 55-64; see also col. 4, lines 30-48 et seq.; see also drawing Figure 2); and

c) wherein coherency of the data transferred is maintained through the use of a snapshot map (see disclosure that copies of blocks on the mass storage system are placed in preservation memory whenever they are going to be changed by a write operation, unless an entry for that block is already in the preservation memory, and furthermore that when a read is requested, the preservation memory is first checked to see if it contains a copy of the block, and that if the copy exists in preservation memory, that copy is returned, otherwise the block is read from the mass storage system, col. 2,

lines 55-64; see also col. 4, lines 30-48 et seq.; see also drawing Figure 2; see also disclosure of block association memory [i.e. the snapshot map] that is used to associate blocks stored in preservation memory with the unique addresses of blocks on the mass storage system, col. 4, lines 51-62).

It would have been obvious to one of ordinary skill in the art at the time of the invention to incorporate snapshot logic of **Ohran et al.** into the storage device controller of **Meyer**, since this would allow the system to create periodic backups for recovery in the event of a failure of the mass storage system, while ensuring that said periodic backup would not be rendered inconsistent in the case where said mass storage system was being updated by other programs as the backup copy is being made, col. 1, lines 20-50).

13. Regarding claim 40, **Meyer** teaches a system substantially as claimed, comprising:

- a) a storage device controller that is operable to receive a copy command specifying the source volume and a target volume (see disclosure that the data storage controller transfers data directly between the first and second

data storage devices under control of the data storage device controller, without employing the memory array and computer bus, said transferring constituting copying, rendering the claimed copy command specifying the source volume and a target volume inherent, col. 4, lines 28-44; see also col. 2, lines 7-19);

- b) the controller being operable to receive a write command specifying the source volume (see col. 5, lines 48-60); and
- c) the controller being operable in response to receiving the copy command to generate and send one or more storage device commands to one or more storage devices for the source and target volumes to copy data from the source volume directly to the target volume without having a file server in the data path (see disclosure that the data storage controller transfers data directly between the first and second data storage devices under control of the data storage device controller, without employing the memory array and computer bus, said transferring constituting copying, rendering the claimed copy logic inherent, col. 4, lines 28-44; see also col. 2, lines 7-19).

Meyer does not explicitly teach a system including snapshot logic.

Ohran et al., however, teaches a system for providing and maintaining snapshots, including:

- a) a replication manager operable to issue a snapshot command (see Abstract, disclosing that the reference is a method for providing a static snapshot; see also col. 1, lines 15-18; see also disclosure of a user indicating that a static image of the mass storage system is desired, said indication being analogous to the claimed snapshot command, col. 4, lines 14-24);
- c) the system being operable to communicate with a replication manager to receive a snapshot command issued by the replication manager, the snapshot command specifying a range of data bytes of a source volume (see disclosure of a user indicating that a static image of the mass storage system is desired, said indication being analogous to the claimed snapshot command, col. 4, lines 14-24; note also the disclosure that mass storage system 104 can be any writable block-addressable storage system, such as one or more disks or a partition of a disk, a partition being a fixed portion of a disk, col. 3, lines 50-56; see also col. 5, lines 23-41);
- d) the system being operable, in response to the snapshot command, to take a snapshot of the range, the snapshot including a snapshot map and snapshot data, the snapshot map being stored by the snapshot logic in the

internal cache and the snapshot data being stored by the snapshot logic in a snapshot volume (see col. 4, lines 20-35; see also disclosure that preservation memory [i.e. the snapshot data] can be an area of memory, one or more disks, a partition of a disk, or a file stored on a disk, col. 3, line 66 through col. 4, line 1; see also disclosure of block association memory [i.e. the snapshot map] that is used to associate blocks stored in preservation memory with the unique addresses of blocks on the mass storage system, col. 4, lines 51-62); and

- e) wherein the snapshot map and snapshot data are used to maintain coherency of any data that is requested (see disclosure that copies of blocks on the mass storage system are placed in preservation memory whenever they are going to be changed by a write operation, unless an entry for that block is already in the preservation memory, and furthermore that when a read is requested, the preservation memory is first checked to see if it contains a copy of the block, and that if the copy exists in preservation memory, that copy is returned, otherwise the block is read from the mass storage system, col. 2, lines 55-64; see also col. 4, lines 30-48 et seq.; see also drawing Figure 2).

It would have been obvious to one of ordinary skill in the art at the time of the invention to incorporate snapshot logic of **Ohran et al.** into the storage device controller of **Meyer**, since this would allow the system to create periodic backups for recovery in the event of a failure of the mass storage system, while ensuring that said periodic backup would not be rendered inconsistent in the case where said mass storage system was being updated by other programs as the backup copy is being made, col. 1, lines 20-50).

14. Regarding claims 25 and 33, **Ohran et al.** additionally teaches a system and method wherein:

a) the range of the storage volume specified by the snapshot command is a first range, and the write command specifies a second range of data bytes of the source volume (see disclosure of a user indicating that a static image of the mass storage system is desired, said indication being analogous to the claimed snapshot command, col. 4, lines 14-24; note also the disclosure that mass storage system 104 can be any writable block-addressable storage system, such as one or more disks or a partition of a disk, a partition being a fixed portion of a disk, col. 3, lines 50-56; see also col. 5, lines 23-41; see

also disclosure of the intercepting of write commands to the source volume, col. 4, lines 35-41); and

b) the controller is operable, in response to receiving the write command while the source volume is being copied to the target volume, to hold the write command in the cache, check if the first range overlaps with the second range and, if so, copy the second range from the source volume to the snapshot volume, update the snapshot map, and then allow the write command to write to the source volume (see disclosure in the Abstract; see detailed disclosure of this process at col. 5, line 48 through col. 6, line 40; see also flowchart illustrated in Figure 2).

15. Regarding claims 26 and 34, **Ohran et al.** additionally teaches a system and method wherein the replication manager is executed on a file server (see col. 6, lines 50-55).

16. Regarding claims 28, 36 and 41, **Ohran et al.** additionally teaches a system and method wherein the replication manager is operable to control multiple storage device controllers (see col. 6, lines 40-49; see additionally the disclosure in **Meyer** that the data

storage device controller includes first and second device controllers coupled to the first and second storage devices, respectively, col. 4, lines 19-21).

17. Regarding claims 29 and 37, **Ohran et al.** additionally teaches a system and method wherein the one or more storage device commands include SCSI commands (see disclosure that the system includes a mass storage device that could be a SCSI device, col. 3, lines 60-65).

18. Regarding claim 45, **Ohran et al.** additionally teaches a system wherein a block size is specified so that fixed size blocks are written to the destination storage device (see col. 5, lines 23-41).

19. Claim 46 is rejected under 35 U.S.C. 103(a) as being unpatentable over **Ohran et al.** (U.S. Patent 5,649,152) in view of **Meyer** (U.S. Patent 5,867,733).

20. Regarding claim 46, **Ohran et al.** teaches a method substantially as claimed, comprising:

- a) receiving a write request, wherein the write request has a range of data bytes to modify (see disclosure that all write commands are intercepted, col. 4, lines 25-48);
- b) determining the range of bytes is within a snapshot range, wherein a snapshot is a process of copying data to a new location before the data is modified by a write operation (see disclosure of the desire to create a static image of a mass storage system, col. 4, lines 14-19; see also col. 1, lines 20-26; see also col. 2, lines 49-51);
- c) determining the range of bytes has not been snapshotted (see disclosure that copies of blocks on the mass storage system are placed in a preservation memory whenever they are going to be changed by a write operation, unless an entry for that block is already in the preservation memory, col. 2, lines 55-58);
- d) in response to determining the range of bytes has not been snapshotted, copying the range of bytes from a source volume to a snapshot volume (see disclosure that copies of blocks on the mass storage system are placed in a preservation memory [snapshot volume] whenever they are going to be changed by a write operation, unless an entry for that block is already in the preservation memory, col. 2, lines 55-58; see also col. 5, lines 50-53);

- e) updating a volume snapshot map, wherein the snapshot map determines which blocks are located in the snapshot volume, in response to copying the range of bytes (see disclosure that block association memory [the snapshot map] is used to associate blocks stored in the preservation memory [the snapshot volume] with the unique addresses of blocks on mass storage system [the source volume], col. 4, lines 56-62; see also disclosure that the block association memory [snapshot map] is updated as necessary when data is placed within the preservation memory [snapshot volume], col. 5, lines 50-53);
- f) modifying the range of bytes of data from the write request, in response to copying the range of bytes (see disclosure that the data to be written by the mass storage write operation is written to the location on the mass storage system specified by the mass storage write address, col. 5, lines 61-65); and
- g) copying the range of bytes from the source volume to a target volume using the volume snapshot map and data stored in the snapshot volume (see disclosure of the creation of a backup copy/static image of the mass storage system, col. 1, lines 20-30; see also col. 2, lines 21-30 and 49-51; see also col. 4, lines 14-19; see also disclosure that during a read of the virtual device, the preservation memory [snapshot volume] is first checked to see if it

contains a copy of the block from the specified location, and if so, that copy is returned, and if not, the block is read from the mass storage system, col. 2, lines 58-64).

Ohran et al. does not explicitly teach a method wherein the copying is performed without having a file server in the path.

Meyer, however, teaches a method of copying data wherein the copying is performed without having a file server in the path (see disclosure that the data storage controller transfers data directly between the first and second data storage devices under control of the data storage device controller, without employing the memory array and computer bus, said transferring constituting copying, rendering the claimed copy logic inherent, col. 4, lines 28-44; see also col. 2, lines 7-19).

It would have been obvious to one of ordinary skill in the art at the time of the invention to transfer data directly from one device driver to another without having a file system in the path, since this would allow for a direct movement of blocks of data between storage devices without processor intervention, and without using I/O or processor bus bandwidth (see col. 4, lines 45-57).

21. Claims 24, 27, 32 and 35 are rejected under 35 U.S.C. 103(a) as being unpatentable over **Meyer** (U.S. Patent 5,867,733) in view of **Ohran et al.** (U.S. Patent 5,649,152) as applied to claims 23, 25, 26, 28, 29, 31, 33, 34, 36, 37, 39-41 and 45 above, and further in view of **Tawil** (U.S. Patent 6,421,723).

22. Regarding claims 24 and 32, **Meyer** and **Ohran et al.** teach a storage device controller and method substantially as claimed.

Neither **Meyer** nor **Ohran et al.** explicitly teaches a storage device controller and method wherein the storage device is a RAID controller.

Tawil, however, teaches the use of a conventional RAID controller (see col. 3, lines 63-67; see also col. 4, lines 1-11).

It would have been obvious to one of ordinary skill in the art at the time of the invention to incorporate a RAID array to the system of **Meyer** and **Ohran et al.**, since it

is well known in the art that the use of RAID arrays provides redundancy which prevents data loss in the event of a data storage device failure.

23. Regarding claims 27 and 35, **Meyer and Ohran et al.** teach a storage device controller and method substantially as claimed.

Neither **Meyer** nor **Ohran et al.** explicitly teaches a storage device controller and method wherein the file server is connected to a storage area network switch and the file server communicates with the storage device controller through the storage area network switch.

Tawil, however, teaches the use of a storage area network (see col. 1, lines 30-42).

It would have been obvious to one of ordinary skill in the art at the time of the invention to incorporate a storage area network, since they offer centralized storage of data for increased efficiency and data handling, and provide data access reliability and availability, unobtrusive capacity expansion, improved data backup and recovery, and performance that is competitive with local data storage (see col. 1, lines 30-36).

24. Claims 30 and 38 are rejected under 35 U.S.C. 103(a) as being unpatentable over **Meyer** (U.S. Patent 5,867,733) in view of **Ohran et al.** (U.S. Patent 5,649,152) as applied to claims 23, 25, 26, 28, 29, 31, 33, 34, 36, 37, 39-41 and 45 above, and further in view of **Dulai et al.** (U.S. Patent 6,205,479).¹

25. Regarding claims 30 and 38, **Meyer** and **Ohran et al.** teach a storage device controller and method substantially as claimed.

Neither **Meyer** nor **Ohran et al.** explicitly teaches a storage device controller and method wherein the controller is operable to send the one or more storage device commands by using one of an in-band protocol or an out-of-band protocol.

Dulai et al., however, teaches a storage device controller and method wherein the controller is operable to send the one or more storage device commands by using one of an in-band protocol or an out-of-band protocol (see disclosure of the use of an in-band protocol, claims 18 and 21).

It would have been obvious to one of ordinary skill in the art at the time of the invention to utilize an in-band protocol, since this allows the transmission of commands over a widely dispersed network where the use of an out-of-band protocol might be impractical.

26. Claims 42-44 are rejected under 35 U.S.C. 103(a) as being unpatentable over **Meyer** (U.S. Patent 5,867,733) in view of **Ohran et al.** (U.S. Patent 5,649,152) as applied to claims 23, 25, 26, 28, 29, 31, 33, 34, 36, 37, 39-41 and 45 above, and further in view of **Simpson et al.** (U.S. Patent 6,128,306).

27. Regarding claims 42-44, **Meyer** and **Ohran et al.** teach a storage device controller and method substantially as claimed.

Neither **Meyer** nor **Ohran et al.** explicitly teaches a storage device controller and method comprising a list of blocks to be copied which is reordered to optimize copy speed, wherein control data is inserted before and after the source data block, nor wherein the list is buffered.

Simpson et al., however, teaches a storage device controller and method comprising a list of blocks to be copied which is reordered to optimize copy speed (see col. 2, lines 15-18), wherein control data is inserted before and after the source data block (see col. 2, lines 5-9), and wherein the list is buffered (see col. 1, lines 55-58).

It would have been obvious to one of ordinary skill in the art at the time of the invention to include prioritized buffering of output data, since this allows more flexible handling of outgoing data traffic, and furthermore since input/output buffering and prioritization and reordering of data in queues was well known in the art at the time of the invention.

Response to Arguments

28. Applicant's arguments filed 9 December 2006 have been fully considered but they are not persuasive.

29. Regarding the Applicants' argument that the **Meyer** reference teaches away from the Applicants' invention by teaching the direct transfer of data between devices, the

examiner respectfully points out that the rejection of record is based upon the *combination* of the **Meyer** and **Ohran et al.** references.

One cannot show nonobviousness by attacking references individually where the rejections are based on combinations of references. See *In re Keller*, 642 F.2d 413, 208 USPQ 871 (CCPA 1981); *In re Merck & Co.*, 800 F.2d 1091, 231 USPQ 375 (Fed. Cir. 1986).

30. Regarding the Applicants' argument that the **Ohran et al.** reference does not disclose copying data from a source volume to a target volume using a snapshot map to maintain coherency, once again the examiner points out that the examiner respectfully points out that the rejection of record is based upon the *combination* of the **Meyer** and **Ohran et al.** references, and that one cannot show nonobviousness by attacking references individually where the rejections are based on combinations of references.

Furthermore, the **Ohran et al.** reference does in fact disclose copying data from a source volume to a target volume using a snapshot map to maintain coherency. It is disclosed at col. 1, lines 20-50 that the entire purpose of the disclosed invention is to allow the creation of a coherent backup/static image of a mass storage device while allowing the data stored thereon to remain in continuous service. See also col. 2, lines 21-30.

The creation of the static image of the mass storage device is analogous to the claimed copying of data from the source volume to the target volume, and **Ohran et al.** clearly discloses that this process is accomplished through the use of the block association memory [snapshot map] and preservation memory [snapshot volume], as cited in the rejections of record. See particularly the rejection of new claim 46.

31. The rejections of record are maintained by the examiner.

Conclusion

Any inquiry concerning this communication or earlier communications from the examiner should be directed to Luke S. Wassum whose telephone number is 571-272-4119. The examiner can normally be reached on Monday-Friday 8:30-5:30, alternate Fridays off.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, John R. Cottingham can be reached on 571-272-7079. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

In addition, INFORMAL or DRAFT communications may be faxed directly to the examiner at 571-273-4119. Such communications must be clearly marked as INFORMAL, DRAFT or UNOFFICIAL.

Customer Service for Tech Center 2100 can be reached during regular business hours at (571) 272-2100, or fax (571) 273-2100.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.



Luke S. Wassum
Primary Examiner
Art Unit 2167

lsw
7 August 2007